

US-PAT-NO: 6374332

DOCUMENT-IDENTIFIER: US 6374332 B1

**TITLE: Cache control system for performing multiple
outstanding
 ownership requests**

DATE-ISSUED: April 16, 2002

US-CL-CURRENT: 711/145, 711/150 , 711/151

APPL-NO: 09/ 409756

DATE FILED: September 30, 1999

PARENT-CASE:

**CROSS-REFERENCE TO OTHER APPLICATIONS AND ISSUED
PATENT**

**The following co-pending applications of common assignee
contain some common
disclosure:**

**"A Directory-Based Cache Coherency System", filed Nov. 5,
1997, Ser. No.
08/965,004, incorporated herein by reference in its entirety;**

**"Message Flow Protocol for Avoiding Deadlocks", U.S. Pat. No.
6,014,709,
issued Jan. 11, 2001, incorporated herein by reference in its
entirety;**

"High-Speed Memory Storage Unit for a Multiprocessor System Having Integrated Directory and Data Storage Subsystems", filed Dec. 31, 1997, Ser. No. 09/001,588, incorporated herein by reference in its entirety; and

"Directory-Based Cache Coherency System Supporting Multiple Instruction Processor and Input/Output Caches", filed Dec. 31, 1997, Ser. No. 09/001,598, incorporated herein by reference in its entirety; and

"Directory-Based Cache Coherency System Supporting Multiple Instruction Processor and Input/Output Caches", a Divisional of Ser. No. 09/001,598, filed Aug. 24, 2000, Ser. No. 09/645,233, incorporated herein by reference in its entirety.

----- KWIC -----

**US Patent No. - PN (1):
6374332**

Detailed Description Text - DETX (35):

From the above discussion, it is apparent that if a large number of requests are being processed across the MI Interfaces, the necessity to request exclusive ownership from the MSU may substantially increase the

time required
to perform a write operation. The current invention minimizes the
time
required to obtain exclusive ownership by prefetching ownership
before a write
request is actually being processed.

Detailed Description Text - DETX (36):

**Description of the Ownership Prefetching System of the Current
Invention**

Detailed Description Text - DETX (47):

To prevent IP "stalls" from occurring during write requests, an
ownership
prefetch mechanism is implemented which minimizes the delay in
obtaining
ownership of a cache line that is not present within an IP's SLC.
When the
Processing Logic 402 is writing a modified operand to memory, the
requested
write address is presented to the FLC 404. If a cache hit occurs,
the write
operation occurs to the FLC. Regardless of whether a cache hit
occurs to the
FLC 404, the updated data will also be written to SLC.

Detailed Description Text - DETX (67):

It may be noted that a mechanism similar to that provided by the
current
invention for write requests could likewise be implemented for
read requests.
That is, a system for providing multiple read requests for cache
lines not

present in the SLC could be implemented in a manner similar to that shown in FIG. 4 for write requests. However, a design choice was made to exclude this logic for read requests in the preferred embodiment of the current system for several reasons. First, a large percentage of read operations involve instruction fetches. During the execution of a sequence of instructions, instruction execution is often re-directed by the occurrence of a jump, skip, or other such instruction. Obtaining a read copy of a cache line that is subsequently determined to be unneeded because execution re-direction has occurred can waste system resources. Thus, for many read situations, it is considered undesirable to obtain a prefetched copy of the read data. Additionally, since a cache line including a block of instructions should not, in most instances, undergo modification, it will not be exclusively owned by any cache in the system. Thus, even if the MSU does not own a requested cache line, only read access has been provided by the MSU to other caches in the system. As a result, the MSU need not initiate a return operation to fetch ownership and/or updated data, and a request for the cache line may be processed without delay. Thus, a shorter access time is generally associated with many read requests as compared to the time required to complete the

**average write request, making it less necessary to bury the read
access times
following a read miss to an SLC 360.**